CLAIMS

What is claimed is:

[Note: Bracketed **bold *and italicized* cross-referencing text** is provided in the below claims as an aid for readability and for finding corresponding (but not limiting) examples of support in the specification. The bracketed text is not intended to add any limitation whatsoever to the claims and should be deleted in all legal interpretations of the claims and should also be deleted from the final published version of the claims.]

1. A Programmably-Sliceable Switch-fabric Unit (PSSU) having a capability of functioning as an NxN' crossbar, and also having a capability of being programmably sliced to instead function as a plurality of SxS' virtual switch slices, where S<N and S'<N', wherein N does not have to equal N', and S does not have to equal S',

said PSSU **[*207*]** comprising:

(a) absolute Ingress ports (a$\mathbb{I}$'s) **[*111*]** and absolute Egress ports (aE's) **[*119*]** that are alternatively identifiable as Relative ingress ports (Ri's) and Relative egress ports (Re's) of respective ones of said virtual switch slices; where the a$\mathbb{I}$'s can receive routing requests and payloads that are to be routed, and where the aE's can output routed payloads to corresponding destinations; and

(b) a request translator **[*260*]**, operatively coupled to the a$\mathbb{I}$'s for receiving routing requests from the a$\mathbb{I}$'s, said request translator having:

(b.1) first means **[*311*]** for determining, based on the absolute Ingress port identification ($aI_x$) of an ingress port on which a given request arrived, what virtual switch slice a corresponding payload signal belongs to;

(b.2) second means **[*312*]** for determining from a Relative egress port (or ports) identification or identifications (Re's) specified directly or indirectly **[*805*]** in the given request, and from the identification **[*SLICE#*]** of the virtual switch slice provided by said first means, which absolute Egress port (or ports, aE's) the corresponding payload signal is to egress from in accordance with the given request; and

(b.3) third means **[*315,826-827*]** for altering the given request before the request is submitted to a scheduler so that the altered request asks for egress of the corresponding payload signal from said determined aE or aE's.

2. A machine-implemented method **[*300*]** that is carried out in a Programmably-Sliceable Switch-fabric Unit (PSSU) **[*207*]** having a capability of functioning as an NxN' crossbar, where said method gives the PSSU a capability of being programmably sliced to instead function as a plurality of SxS' virtual switch slices, where S<N and S'<N', wherein N does not have to equal N', and S does not have to equal S', where the PSSU includes absolute Ingress ports (aI's) **[*111*]** and absolute Egress ports (aE's) **[*119*]** that are alternatively

identifiable as Relative ingress ports (Ri's) and Relative egress ports (Re's) of respective ones of said virtual switch slices; where the aI's can receive routing requests and payloads that are to be routed, and where the aE's can output routed payloads to corresponding destinations; and where said method comprises:

(a) receiving **[*310*]** routing requests from the aI's

(b) first determining **[*311*]**, based on the absolute Ingress port identification ($aI_x$) of an ingress port on which a given request arrived, what virtual switch slice a corresponding payload signal belongs to;

(c) second determining **[*312*]** from a Relative egress port (or ports) identification or identifications (Re's) specified directly or indirectly **[*805*]** in the given request, and from the identification **[*SLICE#*]** of the virtual switch slice provided by said first determining step, which absolute Egress port (aE, or ports, aE's) the corresponding payload signal is to egress from in accordance with the given request; and

(d) altering **[*315,826-827*]** the given request before the request is submitted to a scheduler so that the altered request asks for egress of the corresponding payload signal from said determined aE or aE's.

3. A machine-implemented method for using a Programmably-Sliceable Switch-fabric Unit (PSSU) that has a capability of functioning as an

NxN' crossbar, and also has a capability of being programmably sliced to instead function as a plurality of SxS' virtual switch slices, where S<N and S'<N', and where the PSSU includes absolute Ingress ports (aI's) **[*111*]** and absolute Egress ports (aE's) **[*119*]** that are alternatively identifiable as Relative ingress ports (Ri's) and Relative egress ports (Re's) of respective ones of said virtual switch slices; where the aI's can receive routing requests and payloads that are to be routed, and where the aE's can output routed payloads to corresponding destinations;

said machine-implemented, usage method being carried out when the PSSU is functioning as said plurality of SxS' virtual switch slices, and the method comprising:

(a) determining if an ingress-related error rate exceeding a predefined and corresponding threshold is observed for requests and/or payloads arriving through a given ingress port;

(b) if the exceeding error rate is observed, determining which virtual switch slice is associated with the given ingress port whose ingress-related error rate is exceeding the predefined and corresponding threshold, and disabling that virtual switch slice.

4. A request translation method **[*300*]** for use in a system where absolute Ingress ports (aI's) and absolute Egress ports (aE's) are alternatively identifiable as Relative ingress ports (Ri's) and Relative egress ports (Re's) of

respective, switch slices, and where the translation method comprises:

(a) determining, based on the absolute Ingress port identification ($aI_x$) of the port on which a given request arrived, what slice a corresponding payload signal belongs to;

(b) determining from a Relative egress port (or ports) identification (Re) specified in the given request, which absolute Egress port (or ports, aE's) the corresponding payload signal will egress from; and

(c) altering the given request so as to cause the corresponding payload signal to egress from said determined aE or aE's.

5. A method **[*Figs8A-8E*]** for configuring a Programmably-Sliceable Switch-fabric Unit (PSSU) **[*207*]** having a capability of functioning as an NxN' crossbar, where said PSSU also has a capability of being programmably sliced to instead function as a plurality of SxS' virtual switch slices, where S<N and S'<N', where the PSSU includes absolute Ingress ports (aI's) **[*111*]** and absolute Egress ports (aE's) **[*119*]** that are alternatively identifiable as Relative ingress ports (Ri's) and Relative egress ports (Re's) of respective ones of said virtual switch slices; where the aI's can receive routing requests and payloads that are to be routed, and where the aE's can output routed payloads to corresponding destinations; and where said configuring method comprises:

(a) programming the PSSU **[*Figs8A-8C*]** to be sliced such that its plural

virtual slices are intermingled to provide effective uncrossings of the ingress and/or egress signal lines of the PSSU.

6. A testing method for testing components that are to be used in a production switching system **[*100*]** having a plurality of PSSU's (Programmably-Sliceable Switch-fabric Units) defining a switch fabric **[*105*]**, where the switch fabric is to be operatively connected by way of an interconnect layer **[*103*]** to a lines-interfacing layer **[*101*]**, said components testing method comprising:

(a) providing a test bed system which comes up (e.g., boots up) in a correspondingly simple, programmable slicing mode having 2x2, 4x4, or alike dimensioned, small virtual switch slices, the number of line card units in the test bed system being substantially less than the number of line card units that are to be used in the production switching system **[*100*]**;

(b) testing in the test bed system, the operability of interconnecting components of at least one of the lines-interfacing layer **[*101*]** and the switch fabric **[*105*]** by applying permutated routing test patterns through the relatively small virtual switch slices;

(c) after operability of the switching system components have been verified in the simpler test bed system, moving the tested components to a more complex, production switching system which uses a programmable slicing mode having 8x8, 16x16, or alike dimensioned, larger virtual switch slices.

7. An inventory utilization method **[*Fig1A*]** for servicing a segmented market having an MxM'-ports switching segment **[*40*]** and a SxS'-ports switching segment **[*10*]**, where M and S are different whole numbers each greater than one, and also where M' and S' are corresponding different whole numbers each greater than one, said inventory utilization method comprising:

(a) maintaining a common inventory of Programmably Sliceable Switch Units (PSSU's) that can each be programmably configured to function in a first mode **[*SAP-S*]** as a plurality of K/S, first switching slices with each such first slice providing an S by S' switch matrix capability, and that can each be programmably configured to function in a second mode **[*SAP-M*]** as one or more, second switching slices (K/M slices) with each such second slice providing an M by M' switch matrix capability, where K is greater than S and equal to or greater than M;

(b) in response to demand in the SxS'-ports switching segment:

(b.1) first removing from said common inventory, one or more of the PSSU's,

(b.2) coupling the first removed PSSU's with first software **[*55*]** that configures them to each operate as up to K/S, first switching slices, and

(b.3) supplying the first removed PSSU's with the correspondingly coupled first software to in the SxS'-ports switching segment of the market; and

(c) in response to demand in the MxM'-ports switching segment,

(c.1) second removing from said common inventory one or more of the PSSU's,

(c.2) coupling the second removed PSSU's with second software **[*56*]** that configures them each to operate as up to K/M, second switching slices, and

(c.3) supplying the second removed PSSU's with the correspondingly coupled second software to in the MxM'-ports switching segment of the market.

8. The inventory utilization method **[*Fig1A*]** of Claim 7 and further comprising:

(d) in response to demand slackening in at least one of the SxS'-ports and MxM'-ports switching segments, returning to said common inventory at least some of the PSSU's that had been supplied to the now-slackened, one or more of the SxS'-ports and MxM'-ports switching segments.

9. A monolithically integrated switch-fabric circuit **[*Fig.1B*]** comprising:

(a) a plurality of ingress ports **[*111*]**, where each ingress port is adapted to receive ingressing request signals **[*250a*]** that request switching of a corresponding payload signal, and where each ingress port is further adapted to receive ingressing payloads **[*852c,250b*]** which are to be switched in

accordance with corresponding ones of the request signals;

(b) a plurality of egress ports **[*119*]**, where each egress port is adapted to output switched ones of said ingressing payloads, where identification of which egress port is to output which of the ingressing payloads is provided by egress identification indicia included in said ingressing request signals and the egress identification indicia provides said identification by way of at least one of:

(b.1) a direct and absolute identification ($aE_x$) of a respective one or more of the egress ports that are to output corresponding ones of the switched payloads;

(b.2) a direct and relative identification ($Re_x$) of a respective one or more of the egress ports that are to output corresponding ones of the switched payloads, where the direct and relative identification ($Re_x$) is based on a switch slice **[*VS#0,1,2,3*]** that is assumed to be coupled to the ingress port through which the corresponding request signal ingressed;

(b.3) an indirect, but nonetheless absolute identification ($aE_x$) of a respective subset of the egress ports that are to multicast out a corresponding one of the switched payloads; and

(b.4) an indirect and relative identification **[*580,602,612*]** of a respective subset of the egress ports that are to multicast out a corresponding one of the switched payloads, where the indirect

and relative identification **[*580*]** is based on a switch slice **[*VS#0,1,2,3*]** that is assumed to be coupled to the ingress port through which the corresponding request signal ingressed;

(c) pre-switch, signal processing resources **[*112,114*]** coupled to the plurality of ingress ports for processing at least said ingressing request and payload signals;

(d) post-switch, signal processing resources **[*117,118*]** coupled to the plurality of egress ports for processing at least said switched and egressing payload signals;

(e) a switch matrix **[*116*]** interposed between said pre-switch and post-switch, signal processing resources for switching corresponding ones of ingressing payloads that have been processed by the pre-switch, signal processing resources for egress through the post-switch, signal processing resources and through request-identified ones of said egress ports;

(f) a request translator **[*260*]** operatively coupled to the egress ports for receiving the egress identification indicia **[*262,570,580*]** which represent at least one of said direct and relative identifications ($Re_x$) and said indirect and relative identifications, and for producing translated, egress identification signals **[*596*]** that provide absolute identifications of the egress ports through which egress of corresponding payloads is being requested by the corresponding ingressing request signals.

10. The monolithically integrated switch-fabric circuit **[*Fig.1B*]** of Claim 9 wherein said pre-switch, signal processing resources provide at least four of the following pre-switch, signal processing functions:

(c.1) level-shifting;

(c.2) bit-detection;

(c.3) clock recovery for asynchronously communicated signals;

(c.4) conversion from a relatively serial format to a more parallel data format **[*414*]**;

(c.5) signal framing **[*447*]**;

(c.6) error checking and/or correction **[*150c*]**;

(c.7) ingress-side error logging for identifying ingress links that are currently error-prone **[*Fig3B*]**;

(c.8) signal decoding or decompression;

(c.9) storage of ingress signal data; and

(c.10) alignment of received ingress signals with internal clocks of the switch-fabric circuit.

11. The monolithically integrated switch-fabric circuit **[*Fig.1B*]** of Claim 9 wherein said post-switch, signal processing resources provide at least four of the following post-switch, signal processing functions:

(d.1) encoding and/or compression of egressing payload signals;

(d.2) attachment of error correction **[*150c*]** and other control signals or

framing signals to the egressing payload signals;

(d.3) egress-side error logging for identifying egress links that are currently error-prone;

(d.4) phase-alignment of egressing signals to external output clock signals if synchronous communication is being used;

(d.5) serialization **[*414*]** of egressing signals;

(d.6) level-shifting of egressing signals; and

(d.7) transmission of grant **[*432a*]** and/or protocol signals as may be appropriate in accordance with protocols used by systems with which the switch-fabric circuit communicates.

12. The monolithically integrated switch-fabric circuit **[*Fig.1B*]** of Claim 9 wherein said request translator comprises:

(f.1) a resource parser **[*552*]** for determining whether received request signals are asking for multicast output of corresponding payloads or unicast output of corresponding payloads and for providing a first mode signal **[*560*]** indicating whether a correspondingly-parsed request signal is of the unicast type or multicast type;

(f.2) a multicast lookup table **[*582*]** for converting multicast egress codes into corresponding, specific multicast identifications **[*586,602*]** of specific, relative egress ports from which a corresponding payload is to be output;

(f.3) a multicast mode translation block **[*582*]**, coupled to the multicast

lookup table, for converting the specific multicast identifications of said specific, relative egress ports into absolute multicast identifications; and

(f.4) a unicast mode translation block **[*574*]**, coupled to the resource parser, for converting unicast type, relative relative identifications ($Re_x$) into absolute egress identifications **[*$aE_y$*]**.

13. The monolithically integrated switch-fabric circuit **[*Fig.1B*]** of Claim 12 wherein said request translator further comprises:

(f.5) a slicing pattern register **[*704*]** for indicating what pattern of virtual switch slices is being implemented in the switch matrix; and

(f.6) a source port indicator **[*708*]** for indicating what absolute ingress port ($aI_x$) a corresponding request signal came from;

wherein each of the multicast mode translation block and the unicast mode translation block is coupled to, and responsive to outputs of the slicing pattern register and of the source port indicator.

14. The monolithically integrated switch-fabric circuit **[*Fig.1B*]** of Claim 13 wherein said request translator further comprises:

(f.7) an on-the-fly controlled multiplexer **[*594*]**, coupled to outputs of the multicast mode translation block and of the unicast mode translation block, and driven by said first mode signal **[*560,562*]** for selectively outputting on-the-fly, at least one of the outputs of the multicast mode translation block and of the unicast

mode translation block.

15. The monolithically integrated switch-fabric circuit **[*Fig.1B*]** of Claim 13 wherein said multicast mode translation block performs a variable shift operation **[*Fig6A*]** dependant on said outputs of the slicing pattern register and of the source port indicator.

16. The monolithically integrated switch-fabric circuit **[*Fig.1B*]** of Claim 13 wherein said multicast mode translation block selects amongst a plurality of spread-and-shift operations **[*Fig6B,8E*]** dependant on said output of the slicing pattern register.